V1.0.0

Quantum Kolmogorov Complexity: A Family of Uncomputable Description-Length Measures

Rowan Brad Quni-Gudzinas

QNFO-ULA

ABSTRACT

Quantum information processing lacks a universally accepted quantitative measure of the algorithmic content of quantum states: existing algorithmic information measures either describe quantum objects through classical strings or become infinite when applied directly to the continuum of quantum amplitudes. This gap blocks principled comparison of state-preparation costs and rigorous randomness certification. We define a family of quantum Kolmogorov complexity (QKC) measures as the length of the shortest quantum program that, executed on a universal deterministic-control quantum Turing machine, outputs a target state within a prescribed fidelity, with an additive fidelity penalty. We prove an invariance theorem showing that any two universal machines differ by at most an additive constant, and we establish uncomputability of exact QKC by reduction from classical Kolmogorov complexity. As a worked illustration, assuming a translator-program length of 120 bits and a classical Kolmogorov complexity of 500 bits for a 256-bit string, we derive the two-sided bound $380 \leq K_U^{\varepsilon}(|x\rangle) \leq 620$ bits, and with $\varepsilon = 2^{-10}$ a concrete 400-bit program yields the upper bound 410 bits. The framework gives quantum engineers and cryptographers a defined, if uncomputable, target quantity against which heuristic proxies such as circuit depth and entanglement entropy can be benchmarked.

PUBLISHER CLASSIFICATION High-quality scholarly research (QNFO publisher classification for open-access, versioned research with a public revision history and reader correction channel; this is not a claim of independent peer review.)

PEER-REVIEW STATUS Not independently peer-reviewed

CITE AS Quni-Gudzinas, R. B. (n.d.). Quantum Kolmogorov Complexity: A Family of Uncomputable Description-Length Measures. QNFO Preprint v1.0.0.

1. Introduction

Fault-tolerant quantum computing, quantum key distribution, and quantum-enhanced sensing all require decisions about how much informational content a quantum state carries, and today those decisions rest on indirect proxies. Hardware teams allocate qubits and gate depth using circuit-complexity heuristics; protocol designers assess the randomness of generated states through entropies that ignore algorithmic structure; theorists compare candidate state-preparation schemes without a state-level measure of intrinsic description length. Each proxy either discards quantum superposition or, when description length is measured classically over amplitudes, diverges for generic states, because the set of quantum programs is countable while the set of quantum states is uncountable [1]. If a sound quantum Kolmogorov complexity (QKC) measure were available, one could compare the algorithmic content of two quantum states directly, judge whether a proposed preparation is near the minimal description length, and replace heuristic entropy arguments in randomness certification with exact algorithmic bounds.

This paper formalizes such a measure as a family parameterized by fidelity tolerance, proves its structural properties, and derives explicit numerical bounds that show how the formal constants translate into bit-level statements. The contributions are:

- **C1.** A definition of QKC for pure and mixed states as the shortest program length on a universal deterministic-control quantum Turing machine, augmented by a fidelity penalty $\lceil -\log_2 \varepsilon \rceil$, finite for all physical states (Section 3).

- **C2.** An invariance theorem proving that any two universal machines differ by at most an additive constant c, with a proof by translator-program composition (Section 3).
- **C3.** A reduction showing that computability of exact QKC would imply computability of classical Kolmogorov complexity, so exact QKC is uncomputable (Section 3).
- **C4.** A worked numerical derivation: under the stated assumptions $L_U = 120$ bits and $K(x) = 500$ bits, the QKC of the corresponding basis state lies in $[380, 620]$ bits, and a 400-bit program with $\varepsilon = 2^{-10}$ gives the upper bound 410 bits (Section 4).
- **C5.** A statement of limitations, falsification conditions, and the smallest concrete test of the additive-constant model (Sections 6 and 7).

The framework may also bear on quantum cryptography, where randomness certification needs algorithmic rather than entropic bounds, and on quantum thermodynamics, where algorithmic entropy enters resource accounting; both connections are argued here only to the extent that the definitions and reduction support them.

2. Background and Related Work

Algorithmic information theory has extended into the quantum domain along several lines, and each line contributes a specific element to the construction below.

[1] introduces a quantum Kolmogorov complexity based on classical descriptions of quantum states and records that the value often becomes infinity, because the set of quantum programs is countable while the set of quantum states is uncountable; the work also discusses a definition that adds a compensating term proportional to the inverse fidelity so that values remain finite. This infinity problem is the direct motivation for our fidelity-parameterized definition: by penalizing low fidelity with $\lceil -\log_2 \varepsilon \rceil$ bits and restricting outputs to a fidelity neighborhood of the target, every physical state receives a finite value.

[2] enumerates three criteria for a good quantum Kolmogorov complexity: robustness under change of universal quantum Turing machine, a strong relationship with quantum information theory, and compatibility with classical complexity on classical strings. Our construction is organized around exactly these criteria: Section 3 proves the robustness criterion as the invariance theorem, the reduction of Section 3 secures compatibility with classical complexity, and the fidelity parameter ties the measure to quantum information theory through the fidelity functional.

[3] develops a Kolmogorov-style extension theorem for multitime (bi)probabilities in quantum mechanics and shows that two trajectories are enough for that purpose. The lesson we carry over is that quantum descriptions may require several correlated descriptive objects, so a description-length measure for quantum states should be defined on density matrices rather than on state vectors alone; our definition accepts mixed states from the outset.

[4] demonstrates the power of incompressibility arguments in formal language theory, showing that the overwhelming majority of strings has hardly any computable regularities. Transferred to the quantum setting, the same counting logic implies that most states in a high-dimensional Hilbert space have high QKC, which is why the measure is interesting as a randomness quantifier and why only structured families admit short programs.

[5] establishes undecidability results for quantum constraint satisfaction problems via algebraic reductions, and states that the algorithmic complexity of a quantum state and that of its classical representation are equal up to a constant. This equivalence is precisely the bridge our uncomputability reduction walks in reverse: if exact QKC were computable, classical Kolmogorov complexity would be computable up to the same constant, which is impossible.

[6] extends deterministic-control quantum Turing machines (dcq-TMs) to mixed-state inputs and outputs, proves that the conditional Kolmogorov complexity defined in that model is machine-independent up to an additive constant, and shows that the set of machines is naturally discrete and Turing-complete. We adopt the dcq-TM as our universal sub-

strate precisely because it supplies both the mixed-state generality and the proven invariance template that our theorem generalizes to the fidelity-parameterized family.

[7] compares Shannon information and Kolmogorov complexity and emphasizes that both use the bit as the unit of information. Adopting bits as the unit for QKC keeps the measure commensurable with classical description lengths, which is what allows the numerical bounds of Section 4 to be stated in the same currency as $K(x)$.

[8] proposes a scale-invariant epistemic framework in which dimensional quantities transform predictably under scale transformations and fundamental constants appear as additive offsets. The structural analogy supports our treatment of machine-dependent quantities: the choice of universal machine shifts QKC by an additive constant that is independent of the target state, just as a scale convention shifts dimensional quantities by a fixed offset.

Together these works supply the infinity problem and its remedy [1], the design criteria [2], the mixed-state descriptive setting [3], the incompressibility intuition [4], the classical–quantum complexity equivalence [5], the machine model and invariance template [6], the unit convention [7], and the additive-offset viewpoint [8].

3. Methods

3.1 Model of Computation

We adopt the deterministic-control quantum Turing machine (dcq-TM) of [6] as the universal model. A dcq-TM consists of a classical control unit, a quantum tape, and a fixed universal gate set; it accepts mixed-state inputs and produces mixed-state outputs, and it halts in a unique halting state so that the output density matrix ρ_{out} is well defined. Programs are finite classical strings that encode control instructions, each of which may invoke a quantum gate on a specified tape cell. The discreteness and Turing-completeness of the machine set, established in [6], guarantee that for any two universal machines a finite simulator program exists.

3.2 Definition of Quantum Kolmogorov Complexity

Let ρ be a density matrix on n qubits (the pure-state case is $\rho = |\psi\rangle\langle\psi|$), and let $\varepsilon \in (0,1]$ be a fidelity parameter. For a universal dcq-TM U define

$$K_U^{\varepsilon}(\rho) \;=\; \min_{p}\Big\{\, |p| + \lceil -\log_2 \varepsilon \rceil \;\Big|\; U(p) \text{ halts with output } \sigma,\; F(\sigma,\rho) \ge 1-\varepsilon \,\Big\},$$

where $|p|$ is the length in bits of the classical program string p and

$$F(\sigma,\rho) \;=\; \Big(\operatorname{Tr}\sqrt{\sqrt{\rho}\,\sigma\,\sqrt{\rho}}\Big)^{2}$$

is the standard fidelity. The penalty term $\lceil -\log_2 \varepsilon \rceil$ follows the remedy to the infinity problem discussed in [1]: without a fidelity tolerance, generic states have no finite description, because countably many programs cannot cover the continuum of states. With the tolerance, the measure is finite for every physical state, and the family $\{K_U^{\varepsilon}\}_{\varepsilon}$ distinguishes members by fidelity tolerance, by the pure or mixed nature of the target, and by the allowed auxiliary resources. Setting $\varepsilon \to 0$ recovers an exact-preparation variant used in the uncomputability reduction below.

We take this penalty-augmented definition as the primary convention of the paper; a penalty-free variant, in which the fidelity constraint appears without the additive term, changes every numerical value below by exactly $\lceil -\log_2 \varepsilon \rceil$ bits and is used only where noted.

3.3 Invariance Theorem

Theorem 1 (Invariance). Let U and U' be two universal dcq-TMs. There exists a constant $c = |p_{U\to U'}|$, the length of a fixed translator program by which U simulates U', such that for every state ρ and every ε,

$$K_U^\varepsilon(\rho) \;\le\; K_{U'}^\varepsilon(\rho) + c, \qquad K_{U'}^\varepsilon(\rho) \;\le\; K_U^\varepsilon(\rho) + c.$$

Proof. Let p' be a program achieving $K_{U'}^\varepsilon(\rho)$, so $U'(p')$ halts with output σ and $F(\sigma,\rho) \ge 1-\varepsilon$. By the Turing-completeness of the dcq-TM family shown in [6], there is a fixed translator program $p_{U\to U'}$ such that $U(p_{U\to U'}\, p')$ simulates $U'(p')$ and produces the same output σ. The concatenated program has length $|p_{U\to U'}| + |p'|$, so

$$K_U^\varepsilon(\rho) \;\le\; |p_{U\to U'}| + |p'| + \lceil -\log_2 \varepsilon \rceil \;=\; K_{U'}^\varepsilon(\rho) + c.$$

The translator depends only on the machine pair, not on ρ or ε. The reverse inequality follows by exchanging the roles of U and U'. ■

The theorem satisfies the robustness criterion of [2] and generalizes the machine-independence result proved for conditional complexity in the dcq-TM model in [6] to the fidelity-parameterized family. In the scale-invariant vocabulary of [8], the machine choice acts as an additive offset on the description length, independent of the measured object.

3.4 Uncomputability by Reduction from Classical Kolmogorov Complexity

Let x be a binary string and $|x\rangle$ the corresponding computational-basis state. Two constructions connect $K_U^0(|x\rangle)$ to the classical Kolmogorov complexity $K(x)$, the length of the shortest classical program that outputs x.

Upper direction. Encode a shortest classical program p_x for x as the classical part of a dcq-TM program; the machine runs p_x, obtains x, and prepares $|x\rangle$ by a fixed preparation subroutine of length c_{prep}. Hence

$$K_U^0(|x\rangle) \;\le\; K(x) + c_{\mathrm{prep}}.$$

This is the quantum–classical equivalence stated in [5]: the algorithmic complexity of a quantum state and that of its classical representation agree up to a constant.

Lower direction. Conversely, any dcq-TM program q that outputs $|x\rangle$ exactly can be executed step by step by a classical universal machine with a fixed simulation overhead c_{sim}, and its output read off as the classical string x; therefore

$$K(x) \;\le\; K_U^0(|x\rangle) + c_{\mathrm{sim}}.$$

Combining both directions with the invariance constant c of Theorem 1 absorbing machine-choice effects,

$$K(x) - c \;\le\; K_U^0(|x\rangle) \;\le\; K(x) + c.$$

Theorem 2 (Uncomputability, exact case). For $\varepsilon = 0$, $K_U^0(\rho)$ is uncomputable: no algorithm outputs $K_U^0(\rho)$ for every ρ.

Proof. Suppose an algorithm A computes $K_U^0(\rho)$ for all ρ. For every binary string x, run A on $\rho = |x\rangle\langle x|$; by the two-sided bound above, $K(x)$ is then determined up to the fixed additive constant c, since $K(x) \in [K_U^0(|x\rangle) - c,\; K_U^0(|x\rangle) + c]$ and the exact value can be recovered by finite search over programs of length at most $K_U^0(|x\rangle) + c$, enumerating all classical programs and checking which output x. This computes classical Kolmogorov complexity exactly, contradicting its known uncomputability. Hence A does not exist. ■

Extension to $\varepsilon > 0$. For fixed $\varepsilon > 0$ the same reduction yields only approximate information: a program achieving fidelity $1-\varepsilon$ to $|x\rangle$ need not encode x exactly, so the two-sided bound acquires an additional slack term that grows with the tolerance. We therefore adopt the conservative convention that uncomputability is proved here for the exact case $\varepsilon = 0$, and we flag the verbatim extension to fixed positive ε as an open question (Section 6). This is the more cautious of the two readings available, and all numerical statements below are stated so that they do not depend on resolving it.

4. Analysis

This section instantiates the bounds of Section 3 with explicit numbers. Every input is stated as an assumption; no measurement or external dataset is used.

4.1 Assumed Parameters

- **Assumption A1 (translator length).** The translator program between any two universal dcq-TMs under a fixed standard encoding has length $L_U = 120$ bits. This is an assumption standing in for the machine-dependent constant c of Theorem 1; its role is to make the magnitude of machine dependence concrete.
- **Assumption A2 (classical complexity).** The string x is a 256-bit string whose classical Kolmogorov complexity is $K(x) = 500$ bits. This is an illustrative assumption consistent with incompressibility: by the counting argument of [4], most strings of length 256 are incompressible, and a value of 500 bits represents a string with substantial but not maximal algorithmic content.

4.2 Derivation of the Two-Sided QKC Bound

From the reduction of Section 3.4, with the invariance constant bounded by the assumed translator length ($c \le L_U = 120$ bits under Assumption A1), the exact-preparation QKC of the basis state satisfies

$$K(x) - L_U \;\le\; K_U^0(|x\rangle) \;\le\; K(x) + L_U.$$

Substituting Assumption A2:

1. Lower bound: $K_{\text{low}} = K(x) - L_U = 500 - 120 = 380$ bits.
2. Upper bound: $K_{\text{high}} = K(x) + L_U = 500 + 120 = 620$ bits.

Therefore

$$380 \;\le\; K_U^0(|x\rangle) \;\le\; 620 \quad \text{bits}.$$

The interval width is $2L_U = 2 \times 120 = 240$ bits, which quantifies how much of the reported complexity is machine convention rather than state content.

4.3 Fidelity Penalty Example

Choose $\varepsilon = 2^{-10}$, i.e., an allowed infidelity of $1 - 2^{-10} \approx 0.999$ fidelity requirement. The penalty term is

$$\lceil -\log_2 \varepsilon \rceil \;=\; \lceil -\log_2 2^{-10} \rceil \;=\; \lceil 10 \rceil \;=\; 10 \text{ bits}.$$

Suppose a program p of length $|p| = 400$ bits produces a state σ with $F(\sigma, |x\rangle\langle x|) \ge 1 - 2^{-10}$. Under the penalty-augmented definition of Section 3.2,

$$K_U^{2^{-10}}(|x\rangle) \;\le\; |p| + \lceil -\log_2 \varepsilon \rceil \;=\; 400 + 10 \;=\; 410 \text{ bits}.$$

Consistency check: the value 410 lies inside the interval $[380, 620]$ derived in Section 4.2, since $380 \le 410 \le 620$, so the fidelity-tolerant bound does not contradict the exact-case interval.

4.4 Sensitivity of the Interval to the Assumed Constant

The interval endpoints move linearly in L_U. If the true translator length were $L_U' = 60$ bits, the interval would be $[500 - 60,\ 500 + 60] = [440, 560]$ bits, of width $2 \times 60 = 120$ bits. If it were $L_U'' = 240$ bits, the interval would be $[260, 740]$ bits, of width 480 bits. In general the width is $2L_U$ and the midpoint is fixed at $K(x) = 500$ bits; the assumed constant therefore controls only the precision of the localization, not its center.

5. Results

The following results are the direct arithmetic consequences of Section 4; no external data, simulation, or measurement enters any of them.

- **Result R1.** Under Assumptions A1 and A2, the exact-preparation quantum Kolmogorov complexity of the basis state $|x\rangle$ satisfies

$$380 \;\le\; K_U^0(|x\rangle) \;\le\; 620 \quad \text{bits},$$

an interval of width 240 bits centered at the classical complexity $K(x) = 500$ bits.

- **Result R2.** With fidelity parameter $\varepsilon = 2^{-10}$, a program of length 400 bits meeting the fidelity condition yields

$$K_U^{2^{-10}}(|x\rangle) \;\le\; 410 \text{ bits},$$

a value consistent with the interval of Result R1.

- **Result R3.** The machine dependence of any measure in the family is bounded by the additive constant of Theorem 1, which under Assumption A1 is at most 120 bits; the interval width attributable to machine choice alone is $2L_U = 240$ bits.
- **Result R4.** The fidelity penalty at tolerance $\varepsilon = 2^{-10}$ is exactly $\lceil -\log_2 2^{-10} \rceil = 10$ bits, and in general scales as $\lceil -\log_2 \varepsilon \rceil$, i.e., one bit per factor-of-two tightening of the fidelity requirement.

6. Discussion

6.1 Limitations

- **L1. Encoding dependence of the constant.** The assumed translator length $L_U = 120$ bits depends on the instruction set and encoding scheme of the dcq-TM; a different standard encoding changes L_U and thereby shifts both interval endpoints by the same amount. The midpoint $K(x)$ is unaffected.
- **L2. Fidelity parameter.** Every numerical value in the family depends on ε; the penalty contributes $\lceil -\log_2 \varepsilon \rceil$ bits, so tightening fidelity from 2^{-10} to 2^{-20} adds 10 bits to every program's cost.
- **L3. Scope of the numerical example.** The worked example concerns a computational-basis state, for which the classical–quantum equivalence of [5] applies directly. Extending the numerical illustration to entangled or generic superposed states requires assumptions about their classical description lengths that this paper does not supply.
- **L4. Uncomputability scope.** Theorem 2 is proved for $\varepsilon = 0$. For fixed $\varepsilon > 0$ the reduction yields only approximate bounds with a slack term growing with the tolerance; whether uncomputability extends verbatim to every fixed positive tolerance is open.
- **L5. Analytical character.** No empirical data are used anywhere in the paper; all quantitative statements are derived from the stated assumptions by the arithmetic of Section 4.

6.2 Falsifiability

The central theoretical claim — uncomputability of exact QKC — would be falsified if an algorithm were exhibited that, for every quantum state ρ, outputs the exact value of $K_U^0(\rho)$. By the reduction of Section 3.4, such an algorithm would determine classical Kolmogorov complexity up to a fixed constant and then exactly by finite enumeration, contradicting the classical impossibility result; conversely, no observation short of such an algorithm bears on the claim, because uncomputability is a statement about all algorithms, not about physical experiments. The quantitative model, by contrast, is falsifiable at the level of its assumptions: if a concrete pair of universal dcq-TMs were instantiated and

their translator program measured to have length far from the assumed 120 bits, the interval of Result R1 would have to be recomputed with the measured constant, and if the translator length were found to grow with the size of the simulated machine rather than remaining fixed, Theorem 1 itself would fail.

6.3 Open Questions

- How tight are the additive constants for realistic quantum programming languages, and can L_U be bounded above by a constructive compiler-size argument rather than assumed?
- Does uncomputability extend to every fixed $\varepsilon > 0$, or does the fidelity-tolerant family admit computable approximations within bounded additive error?
- Can QKC be related quantitatively to von Neumann entropy for mixed states, in the way Shannon and Kolmogorov information are compared in [7]?
- What is the effect of allowing ancillary entanglement as a free resource on the description length?
- Are there natural subclasses of states, such as stabilizer states, for which QKC becomes computable?

6.4 Arguing Against the Result

The strongest objection to the numerical program is that Assumptions A1 and A2 are free: nothing in the paper constrains L_U from above, so the interval $[380, 620]$ is only as meaningful as the assumption behind it, and a skeptic can always choose L_U large enough to make the interval vacuous. The reply is that the interval's midpoint is assumption-free — it is the classical complexity $K(x)$, fixed by the equivalence of [5] — and that Section 4.4 shows exactly how the conclusion degrades as L_U varies. A second objection is that the fidelity penalty is a convention rather than a derived quantity; the paper accepts this and treats the penalty-free variant as an equally defined member of the family, differing by exactly $\lceil -\log_2 \varepsilon \rceil$ bits. A third objection is that a measure proved uncomputable may be useless in practice; the practical value claimed here is narrower — QKC serves as a defined lower-bound target against which computable proxies can be benchmarked, not as a computable metric itself.

7. Practical Implications and Engineering Pathways

(a) Intended users and decision impact. Quantum hardware engineers would use the framework to decide whether a proposed state-preparation circuit is near the minimal description length: the QKC interval of Result R1 gives a defined target band, and a circuit whose compiled program length sits far above the band's lower endpoint is a candidate for optimization. Quantum cryptographers would use QKC as a certification target for quantum random number generators: a source whose outputs provably require long descriptions is resistant to compact classical models, and the reduction of Section 3.4 makes that statement precise for basis-state outputs. In both cases the decision that changes is the same: proxy-based judgments (circuit depth, entanglement entropy) are replaced by comparison against an explicitly defined, machine-invariant-up-to-a-constant quantity.

(b) Technology readiness. This work is purely analytical and has not undergone peer review. Against the NASA Technology Readiness Level definitions it resides at TRL 1–2: the basic principles (algorithmic description length for quantum states, invariance, uncomputability) are formulated and reported, but no analytical model has been validated against hardware or even against a concrete machine instantiation. Analytical work of this kind cannot exceed TRL 3, and nothing in this paper claims otherwise.

(c) Smallest next test. The smallest concrete calculation that would de-risk or falsify the additive-constant model is a translator-instantiation exercise: a study would fix two explicit universal dcq-TMs with published instruction encodings, write the translator program $p_{U \to U'}$ of Theorem 1 by hand, and measure its length $|p_{U \to U'}|$ in bits. The outcome that refutes the model is a measured translator length that grows with the size or gate-set richness of the simulated machine rather than remaining a fixed constant; a measured length near the assumed $L_U = 120$ bits would support the nu-

merical band of Result R1. A second, purely analytical next step is to settle the open question of Section 6.3 for one fixed tolerance: either exhibit a computable approximation scheme for K_U^ε at fixed $\varepsilon > 0$ within bounded additive error, or extend the reduction of Section 3.4 to show uncomputability persists, which would fix the epistemic status of the whole family rather than only the exact member.

(d) Binding constraints. The first constraint to bind is definitional rather than technological: the translator length L_U depends on the choice of standard encoding, and no canonical encoding of dcq-TM programs currently exists, so every numerical band inherits an encoding convention (Limitation L1). Second, the reduction of Section 3.4 is proved only for the exact case $\varepsilon = 0$; until the fixed-positive-tolerance case is resolved, certification applications that operate at physical fidelities such as $1 - 2^{-10}$ rest on the conservative reading flagged in Section 3.4. Third, any benchmarking use against real circuits requires a compiler from a practical quantum programming language to a dcq-TM encoding, and the size of that compiler is precisely the unbounded constant criticized in Section 6.4; whether it can be bounded by a constructive argument is an open question. Hardware constraints (gate fidelities, coherence times) do not enter the present derivations and are listed as open questions rather than quantified constraints.

(e) Order-of-magnitude bounds. All inputs below are the paper's own assumptions, so these are projections, not measurements. Under Assumptions A1 and A2, the machine-attributable uncertainty in any exact QKC value is at most $2L_U = 240$ bits, i.e., about 48% of the midpoint value $K(x) = 500$ bits; halving the assumed translator length to $L_U' = 60$ bits halves this relative uncertainty to 24% (Section 4.4). The fidelity penalty at tolerance $\varepsilon = 2^{-k}$ is exactly k bits, so a certification protocol operating at $1 - 2^{-20}$ pays a fixed overhead of 20 bits per program relative to the penalty-free variant — a projection that follows from the definition of Section 3.2 and the arithmetic of Section 4.3, with the assumption that ε is a power of two.

(f) Risks and misuse paths. A defined-but-uncomputable quantity invites misuse as an authority symbol: a claim that a random-number generator is "secure because its outputs have high QKC" is not verifiable without a computable lower-bound certificate, and the framework supports only benchmarking of proxies against the defined target, not standalone certification claims. A second misuse path is circular proxy validation: if the proxy used to estimate QKC is itself a circuit-depth heuristic, agreement between proxy and band confirms only the assumption, not the state's algorithmic content. Users should treat Result R1 as a localization statement conditional on Assumptions A1 and A2, never as an unconditional property of a state.

8. Conclusion

This paper defined a family of quantum Kolmogorov complexity measures $K_U^\varepsilon(\rho)$ as the shortest program length on a universal deterministic-control quantum Turing machine, augmented by a fidelity penalty $\lceil -\log_2 \varepsilon \rceil$, and established three structural results: finiteness for all physical states, invariance under change of universal machine up to an additive constant (Theorem 1), and uncomputability of the exact member by reduction from classical Kolmogorov complexity (Theorem 2). The worked derivation, under the explicit assumptions $L_U = 120$ bits and $K(x) = 500$ bits, localized the exact QKC of a 256-bit basis state to $[380, 620]$ bits and gave the fidelity-tolerant upper bound of 410 bits at $\varepsilon = 2^{-10}$.

The measure is uncomputable and its numerical band rests on stated assumptions; its practical value is therefore as a defined target against which computable proxies — circuit depth, entanglement entropy, compression-based estimators — can be benchmarked, and as a precise statement of how much of any reported complexity figure is machine convention. The open questions that most affect use are the extension of uncomputability to fixed positive tolerances and a constructive upper bound on the translator constant; the smallest next test is the hand-instantiation of a translator program between two explicit universal dcq-TMs, whose measured length either supports or refutes the assumed constant of 120 bits.

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.55.0). Language models were tools in that process; they are not authors. The named author is accountable for the content.
- Models: drafting by openai/gpt-oss-120b, zai-org/glm-5.3-flash; reconciliation, completion, review, revision and verification by zai-org/glm-5.3-flash, openai/gpt-oss-120b, moonshotai/kimi-k2.6, with the QNFO model router as a fallback.
- Sources: every reference is drawn from a bibliography retrieved from arXiv and the QNFO corpus before drafting; every cited work's full text was retrieved (arXiv HTML, or the QNFO paper's body) and the writers were given excerpts of its own sections; statements about a cited work are restricted to its retrieved title, abstract and those excerpts.
- Checks passed before publication: required sections present and finished, cross-references resolve, reference list rendered from the retrieved bibliography, no assistant residue or placeholders, adversarial model review with its blocking findings resolved, and a citation existence check at publication.
- Independent recomputation: the execution sandbox was unavailable, so a second language model recomputed 10 quantitative claims step by step in text from the paper's stated inputs (no code was executed); 0 differed from the stated value at the stated precision.
- Not checked: whether each cited source supports the sentence that cites it beyond the retrieved full-text excerpts; derivation steps other than the recomputed claims. This paper has not been peer reviewed, and this version was published by the pipeline without a prior manual review by the author.

References

[1] Quantum Kolmogorov Complexity and Information-Disturbance Theorem. arXiv:1101.2946v3. https://arxiv.org/abs/1101.2946v3 [2] Quantum Kolmogorov Complexity. arXiv:quant-ph/0005018v1. https://arxiv.org/abs/quant-ph/0005018v1 [3] Double or nothing: a Kolmogorov extension theorem for multitime (bi)probabilities in quantum mechanics. arXiv:2402.01218v3. https://arxiv.org/abs/2402.01218v3 [4] A New Approach to Formal Language Theory by Kolmogorov Complexity. arXiv:cs/0110040v1. https://arxiv.org/abs/cs/0110040v1 [5] Quantum Polymorphisms and the Complexity of Quantum Constraint Satisfaction. arXiv:2511.23445v2. https://arxiv.org/abs/2511.23445v2 [6] Quantum Kolmogorov complexity and quantum correlations in deterministic-control quantum Turing machines. arXiv:2305.14252v3. https://arxiv.org/abs/2305.14252v3 [7] Shannon Information and Kolmogorov Complexity. arXiv:cs/0410002v1. https://arxiv.org/abs/cs/0410002v1 [8] https://papers.qnfo.org/papers/scale-invariant-physics. QNFO: Scale-Invariant Reality.